REMARKS

Claims 1-31 are pending. Claims 1-31 are rejected. Claims 1-11 were rejected under 35 USC 103(a) as being unpatentable over Self (5,623,644) in view of Booth (6,065,073). Claims 12-31 were rejected under 35 USC 103(a) as being unpatentable over Self in view of Kelly (5,379,440), Booth and Dervin (6,952,766.) The rejections are respectfully traversed, for the reasons cited below.

Claims 1

Claim 1 was rejected under 35 USC 103(a) as being unpatentable over Self in view of Booth. Applicants respectfully disagree with the rejection, but have amended claim 1 to expedite prosecution. Claim 1 now recites "disabling the second cluster comprises disabling polling for a link from the first interconnection controller to the second interconnection controller and ~~flushing~~emptying caches associated with the second cluster." The amendment is supported by several sections of the specification, including page 21, lines 1-9, which describe an example method of disabling a cluster:

> A cluster is disabled at 1003 but not yet unplugged. At 1005, caches associated with the cluster to be removed are flushed. At 1007, routing tables are modified to reflect the removal of a processing cluster. At 1009, a fence bit is written in the configuration space registers. A fence bit may be written by a variety of entities. In one example, the fence bit is written by a service processor or by a JTAG interface associated with a processor. At 1011, the multiple processor cluster can be physical removed. At 1013, the physical layer at the interconnection controller still residing in the system may be maintained in order to allow for a new or replacement cluster of processors to be introduced.

The combination of Self and Booth do not teach or suggest the above features. Self pertains to a communication apparatus for communicating messages between two computing resources, such as processors. (See Abstract of Self.) Booth pertains to a system and method of polling a status register within a physical layer interface to a LAN, such as in a network interface card. (See Abstract of Booth.) Neither Self nor Booth involve disabling a cluster or emptying caches. On page 3 of the Office Action, the examiner asserts that Booth teaches the disabling of polling, but the examiner does not cite any reference in either Self or Booth that teaches the disabling of a cluster. Additionally, Booth does not teach the flushing of caches as part of a

process of disabling a cluster. The examiner asserts on page 7 of the Office Action that col. 11, lines 23-28 of Self disclose the flushing of caches in the manner of claim 1. Col. 11, lines 13-31 recite:

> *There is also an entry for each cache line on the processors in each duplicate cache tag store.* This tag store (e.g., 1252 of FIG. 12b) allows the memory controller to filter cache consistency information that is received from the memory network. The processors can use a snoop cache consistency model, as in the prior art, to ensure the validity of their cache. The processors will broadcast any change in state of a cache line and will snoop the input stream for changes that would effect the information that is contained within its cache by examining the memory address of the accesses. *The consistency of caches between micro-clusters is maintained by broadcasting all read and invalidate requests into the memory network to be routed to other micro-clusters in the system. If neither processor in the micro-cluster has a copy of the cache line, then the information is discarded.* Otherwise, the cache consistency information is presented to the processors in the micro-cluster as in the same way in which it would present the information if received from the other processor within the micro-cluster.

(Emphasis added.)

The above citation does not teach the flushing or emptying of a cache in connection with the disabling of a cluster. Instead, Self seems to teach a mechanism for maintaining the consistency of data in different caches associated with different processors. The mechanism may sometimes involve discarding a single cache line. Such discarding of information is meant to be part of a process of operating an enabled micro-cluster, not as part of a process of disabling a cluster. Furthermore, such discarding of information does not involve emptying or completely flushing a cache, but rather removing a cache line from a cache to make it consistent with other caches. As a result, Self and Booth do not teach or suggest disabling a second cluster, "wherein disabling the second cluster comprises disabling polling for a link from the first interconnection controller to the second interconnection controller and emptying caches associated with the second cluster."

Self and Booth also do not teach or suggest processors connected through a point to point link. On page 4 of the Office Action, the examiner, citing Figure 12a of Self, asserts that "Self discloses that processor within a cluster are connected via a point-to-point communication interface employing point to point communication links." Applicants respectfully submit, however, that the examiner is interpreting the claim recitation, "the first processor is connected to the second processor through a point to point link," too broadly. The term "point to point

link" is well known in the art. By way of example, page 120 of the text Sams Teach Yourself TCP/IP in 24 Hours (2004) describes point-to-point connections as follows:

> ...the two computers at either end of a phone line do not have to compete for the transmission medium with other computers—they have to share it only with each other. *This type of connection is called a point to point connection.*

(Emphasis added.)

Applicants do not intend to limit the scope of any term in the claim to the above example. Applicants provide the example, however, to demonstrate that the recitation "the first processor is connected to the second processor through a point to point link" would not be misunderstood by a person of ordinary skill in the art as reading upon an architecture as in Figure 12a of Self, where the two processors 1201 and 1202 lack any direct link with one another and instead are connected to one another only through point to point interface and memory controller 1205. In view of the foregoing, it is respectfully submitted that independent claim 1 is patentable over the prior art of record.

Claims 12 and 23

Independent claims 12 and 23 were rejected under 35 USC 103(a) as being unpatentable over Self in view of Kelly, Booth and Dervin. Applicants respectfully disagree, but have amended claims 12 and 23 to expedite prosecution. Amended claims 12 and 23 now recite "enabling physical layer communications between the first and second interconnection controllers without enabling link layer communications between the first and second interconnection controllers" and "establishing, after asserting the reset signal and enabling the physical layer communications, a link layer protocol on a connection between the first and second interconnection controllers." The amendment is supported by several sections of the specification, including page 20, lines 16-22, which recite:

> In another example, one or more clusters are running and a new cluster needs to be added to the system. According to various embodiments, reset is asserted and deasserted on the new cluster that needs to be added. If the physical layer for a remote connection to the new cluster is disabled, the physical layer placed into the polling state. The remote connection is unfenced. Cluster IDs are exchanged and the physical layer is enabled. The service processor generates

> information needed to program the new cluster and also the main cluster in a similar manner as noted above. The link layer of the new remote connection is now initialized.

The temporary enabling of only physical layer communications may benefit some embodiments of the present invention. Describing Figure 8, page 18, lines 3-5 of the specification recite "at 803, a fence bit is set to disable link layer communications. Disabling link layer communications allows physical layer communications to be established without interference from data local transmissions."

The examiner, citing col. 5, lines 34-48, col. 7, lines 4-col. 8, line 12 and col. 9, lines 7-18, asserts on pages 11-12 of the Office Action that Devin discloses the establishing of a protocol after the initiation of a restart. The cited passages relate to a node initiating a restart operation in the event of the failure of another node. (See col. 7, lines 12-18 of Devin.) In connection with the restart, a "start protocol" is used. (See col. 8, lines 26-28 of Devin.) The examiner is presumably equating the "start protocol" of Devin with the "link layer protocol" of claim 12.

The "start protocol" of Devin, however, does not teach or suggest enabling physical layer communications without link layer communications and, afterward, establishing a link layer protocol. The start protocol is defined in col. 9, lines 19-23 and lines 34-38, which recite:

> Various start protocols may be utilized consistent with the invention. For example, in the aforementioned OS/400 environment, a start protocol may generally operate by using a dedicated TCP/IP port to the inetd daemon to start a cluster-related job on the node... It will be appreciated that other protocols for starting a node may be used in the alternative. For example, other TCP/IP ports, or a shared file system, may be used to start a job for initiating clustering on a node.

The description only states broadly that the start protocol may take different forms. The description does not suggest a sequence of two different types of communications, one of which relates to the physical layer but not to the link layer. Additionally, Devin suggests that the start protocol may utilize a TCP/IP port. TCP and IP, however, relate to layers above that of the physical and link layers. If the start protocol utilized such ports, it would presumably involve both the physical and link layers. Devin does not seem to provide any other suggestion as to the types of layers that the start protocol involves. Thus, Devin does not teach or suggest "enabling

physical layer communications between the first and second interconnection controllers without enabling link layer communications between the first and second interconnection controllers" and "establishing, after asserting the reset signal and enabling the physical layer communications, a link layer protocol on a connection between the first and second interconnection controllers." In view of the foregoing, it is respectfully submitted that claim 12 and 23 be allowed.

The various dependent claims are respectfully submitted to be patentable over the art of record for at least the same reasons as set forth above with respect to their associated independent claims. Furthermore, these dependent claims recite additional features that when considered in the context of the claimed invention, further patentably distinguish the art of record.

Conclusion

Applicant's Attorney believes that all pending claims are allowable and respectfully requests a Notice of Allowance for this application from the Examiner. If any fees are required in connection with the filing of this response, including any fees required for any required extenstion of time, such fees may be charged to Deposit Account No. 500388 (Order No. NWISP046). Should the Examiner believe that a telephone conference would expedite the prosecution of this application, the undersigned can be reached at the telephone number set out below.

Respectfully submitted,
WEAVER AUSTIN VILLENEUVE & SAMPSON LLP

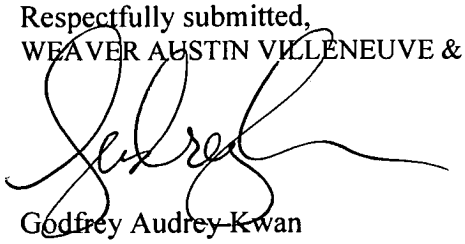

Godfrey Audrey Kwan
Reg. No. 46,850

500 12th Street, Suite 200
Oakland, CA 94607
(510) 663-1100